

Data Center Professional (JNCIP-DC) Exam

Juniper JN0-683

Version Demo

Total Demo Questions: 8

Total Premium Questions: 87

Buy Premium PDF

<https://passqueen.com>

support@passqueen.com

Topic Break Down

Topic	No. of Questions
Topic 1, Data Center Architecture	2
Topic 2, Data Center Deployment	3
Topic 3, Data Center Fabric	72
Topic 4, Data Center Security	3
Topic 5, Data Center Automation and Orchestration	7
Total	87

QUESTION NO: 1

Exhibit.

Referring to the exhibit, which statement is true?

- A. A PBB-EVPN architecture is being used.
- B. An ERB architecture is being used.
- C. An OTT architecture is being used.
- D. A CRB architecture is being used.

ANSWER: B

Explanation:

An ERB architecture is being used. In an Edge-Routed Bridging deployment, the leaf switches act as VXLAN tunnel endpoints and provide the default gateway and Layer 3 routing function for their locally attached VLANs or bridged domains. The configuration indicated in the exhibit associates VLAN-based traffic with VXLAN network identifiers and routing-instance constructs at the leaf. This is the defining operational pattern for ERB: endpoints retain Layer 2 connectivity where required through EVPN-VXLAN, while traffic requiring communication between IP subnets is routed directly at the ingress leaf. As a result, a host sends traffic to its local anycast or distributed gateway, and the leaf performs the Layer 3 lookup before encapsulating traffic across the VXLAN fabric when the destination is remote. EVPN distributes the control-plane information needed for the fabric to identify endpoint reachability and forwarding locations. This model provides distributed routing close to the workload, avoids sending intersubnet traffic to a centralized routing tier, and is a standard Juniper EVPN-VXLAN data-center design. Juniper's EVPN-VXLAN documentation describes the role of VXLAN tunnel endpoints and EVPN control-plane operation in this architecture: [Juniper EVPN-VXLAN Overview](#). Additional Juniper configuration guidance is available in the [EVPN-VXLAN documentation](#).

QUESTION NO: 2

Exhibit.

Referring to the exhibit, which statement is correct?

- A. VNI 100 is not configured on the remote VTEP.
- B. The MAC address is unknown and not in the forwarding table of the remote VTEP.
- C. The remote VTEP is not responding.
- D. The MAC address is known but not reachable by the remote VTEP

ANSWER: B

Explanation:

The MAC address is unknown and not in the forwarding table of the remote VTEP. In an EVPN-VXLAN fabric, a VTEP uses its MAC forwarding information to determine whether a destination MAC address is locally attached, reachable through a remote VTEP, or unknown. If the destination MAC has not been learned or advertised to the remote VTEP, that VTEP does not have an applicable forwarding-table entry for it. This is a normal data-plane condition for an as-yet unlearned destination, rather than evidence that the VTEP itself is unavailable. The VTEP handles traffic for an unknown unicast destination according to the configured VXLAN/EVPN forwarding behavior, which can include flooding within the applicable broadcast domain when ingress replication or multicast replication is used. EVPN distributes MAC/IP reachability information through BGP so that, once the MAC is learned and advertised, participating VTEPs can install the corresponding forwarding state and send subsequent unicast traffic directly to the appropriate VTEP. Juniper's EVPN-VXLAN overview describes this control-plane learning and distribution model: [EVPN-VXLAN Overview](#). Additional Junos EVPN documentation is available at [EVPN-VXLAN Documentation](#).

QUESTION NO: 3

You are implementing seamless stitching between two data centers and have a proposed configuration for a border leaf device.

In this scenario, which two statements are correct? (Choose two.)

- A. The translation-vni must match in both data centers.
- B. The translation-vni must be different in each data center.
- C. The ESI must be different in each data center.
- D. The ESI must match in both data centers.

ANSWER: A C

Explanation:

For EVPN-VXLAN seamless stitching, the translation VNI is an inter-data-center transport identifier. It must be consistent at both ends of the stitched connection so that a border leaf receiving VXLAN-encapsulated traffic can associate the packet with the intended translated bridge domain. Therefore, *The translation-vni must match in both data centers.* is required for the two border domains to stitch the same service correctly.

An Ethernet segment identifier identifies a particular multihomed Ethernet segment in EVPN. Each data center uses its own Ethernet segment for its local border-leaf multihoming and designated-forwarder behavior. Consequently, the Ethernet segments at the separate sites require distinct identifiers, making *The ESI must be different in each data center.* correct. This keeps the sites' local Ethernet-segment state independent while the common translation VNI provides the service stitching across the DCI.

Juniper's EVPN-VXLAN documentation describes EVPN control-plane operation, including Ethernet-segment identification and multihoming behavior, and its data-center-interconnect guidance covers extending EVPN-VXLAN services between fabrics. See [Juniper EVPN-VXLAN Documentation](#) and [EVPN-VXLAN Data Center Interconnect](#).

QUESTION NO: 4

You are designing the EVPN control plane for a large leaf-spine fabric. Every leaf switch is a VTEP, and you want to avoid configuring a full mesh of EVPN BGP sessions between all leaf switches.

Which design would satisfy this requirement?

- A. Use the spine devices as EVPN BGP route reflectors.
- B. Configure a full mesh of EVPN BGP sessions between all leaf devices.
- C. Configure the leaf devices as route reflectors for the spine devices.
- D. Use PIM to distribute EVPN MAC/IP routes between the leaf devices.

ANSWER: A

Explanation:

Use the spine devices as EVPN BGP route reflectors. Each leaf establishes EVPN BGP sessions to the route reflectors rather than to every other leaf. The route reflectors distribute EVPN routes between their clients, which significantly reduces the number of required BGP sessions as the fabric grows.

The route reflectors participate in the EVPN control plane but do not need to be VXLAN tunnel endpoints for tenant traffic. Configuring a direct BGP session between every pair of leaf VTEPs would create the full mesh that the design is intended to avoid.

QUESTION NO: 5

Exhibit.

Both DC and DC2 are using EVPN-VXLAN technology deployed using an ERB architecture. A server on the Red VLAN must communicate with a server on the Green VLAN. The Blue VLAN in DC and DC2 needs to be the same VLAN.

Which statement is correct in this scenario?

- A.** The eight spine devices must be configured as border spine devices; a full mesh interconnect must exist between all eight spine devices and the Blue VLAN must be stitched together
- B.** An interconnect is required between the four SRX Series devices; the Blue VLAN must be stretched and a transit VNI must be added for the Red and Green VLANs.
- C.** An interconnect is required between four leaf devices in the services blocks; the Red VLAN and the Green VLAN must be stitched and the Blue VLAN must be stretched.
- D.** A lean super spine device must be added to DC and DC2; all VLANs must be stretched to the lean super spine device and the lean super spine devices must stitch all the VLANs together.

ANSWER: B

Explanation:

An interconnect is required between the four SRX Series devices; the Blue VLAN must be stretched and a transit VNI must be added for the Red and Green VLANs. This design separates the two required traffic behaviors appropriately. The Blue VLAN has an explicit requirement to represent the same Layer 2 segment in both data centers, so its EVPN/VXLAN service must be extended across the data-center interconnect. EVPN advertises the MAC and IP reachability information needed for endpoints in that stretched VLAN, while VXLAN supplies the overlay encapsulation for transport between sites.

Communication from the Red VLAN to the Green VLAN is inter-VLAN traffic and is routed under an EVPN routed-bridging architecture. The SRX devices provide the secure inter-site connection, and a transit VNI supplies an overlay transit segment through which routed traffic can pass between the sites. This preserves distributed gateway and routing behavior at the EVPN edge while providing a defined, policy-capable DCI path rather than requiring every tenant VLAN to be stretched. Juniper describes EVPN-VXLAN as supporting both Layer 2 bridging and Layer 3 routing services, including centralized or distributed inter-VXLAN routing models. See the [Juniper EVPN-VXLAN overview](#) and [Juniper EVPN-VXLAN routing documentation](#).

QUESTION NO: 6

Exhibit.

```
user@leaf1> show configuration
...
interfaces {
  ge-0/0/0 {
    description "facing_spine1:ge-0/0/1";
    speed 10g;
    mtu 9142;
    unit 0 {
      family inet {
        mtu 9170;
        address 172.16.0.9/31;
      }
    }
  }
  ge-0/0/1 {
    description "facing_spine2:ge-0/0/1";
    speed 10g;
    mtu 9142;
    unit 0 {
      family inet {
        mtu 9170;
        address 172.16.0.11/31;
      }
    }
  }
  irb {
    unit 200 {
      family inet {
        address 192.168.200.1/24;
      }
    }
  }
}
vpls {
  vn100 {
    vpls-id 100;
    description "BLUE";
  }
  vn200 {
    description RED;
    vpls-id 200;
    l3-interface irb.200;
  }
}
```

Host A is connected to vpls 100 on leaf1. Host B is connected to vpls 200 on leaf1. Host A and Host B are unable to communicate. You have reviewed the routing and your hosts have the correct default route (.1)

Referring to the exhibit, which two commands will solve the problem? (Choose two.)

- A. delete vpls vn200 l3-interface irb.200
- B. set interfaces irb unit 100 family inet address 192.168.100.1/24
- C. set routing-options static route 0.0.0.0/0 next-hop 192.168.200.10
- D. set vpls vn100 l3-interface irb.100

ANSWER: B D

Explanation:

The required configuration is to create the Layer 3 IRB gateway for VLAN 100 and associate that IRB unit with VLAN 100. set interfaces irb unit 100 family inet address 192.168.100.1/24 configures the default-gateway address that Host A already expects to use. An IRB interface supplies the routed interface for a bridged VLAN; its unit number is conventionally aligned with the VLAN and is independently assigned an IPv4 address and prefix.

`set vlans vn100 l3-interface irb.100` then binds that newly addressed IRB interface to VLAN 100. With VLAN 100 mapped to `irb.100` and VLAN 200 mapped to its existing `irb.200`, the switch can route traffic between the directly connected 192.168.100.0/24 and 192.168.200.0/24 networks. Because both endpoint default gateways are the local .1 addresses, no additional default static route is required for this inter-VLAN communication; the connected routes created by the two IRB interfaces provide the necessary routing information.

Juniper documents the VLAN `l3-interface` association in the [l3-interface CLI reference](#) and describes IRB operation in its [IRB interfaces overview](#).

QUESTION NO: 7

You are using ingress replication for BUM traffic in an EVPN-VXLAN fabric because multicast routing is not enabled in the underlay.

Which two statements are correct in this scenario? (Choose two.)

- A. The ingress VTEP sends individually encapsulated BUM copies to the applicable remote VTEPs.
- B. The underlay must run PIM to replicate BUM traffic for every VNI.
- C. Ingress replication can operate without multicast-enabled underlay forwarding.
- D. Ingress replication eliminates the need for VXLAN encapsulation.

ANSWER: A C

Explanation:

Ingress replication does not require multicast-enabled underlay forwarding. Instead, the ingress VTEP creates individual VXLAN-encapsulated copies of BUM traffic for the remote VTEPs that participate in the applicable broadcast domain.

This differs from multicast replication, where the underlay multicast network replicates VXLAN traffic toward the remote VTEPs. Ingress replication can simplify the underlay but requires the ingress VTEP to perform the replication work.

QUESTION NO: 8

You are implementing VXLAN broadcast domains in your data center environment. Which two statements are correct in this scenario? (Choose two.)

- A. A VXLAN packet does not contain a VLAN ID.
- B. The VNI must match the VLAN tag to ensure that the remote VTEP can decapsulate VXLAN packets.
- C. Layer 2 frames are encapsulated by the source VTEP.
- D. The VNI is a 16-bit value and can range from 0 through 16.777.215.

ANSWER: A C

Explanation:

“A VXLAN packet does not contain a VLAN ID.” is correct in the VXLAN overlay context because VXLAN identifies the overlay broadcast domain with its VXLAN Network Identifier (VNI), rather than extending a customer VLAN identifier across the IP underlay. A VLAN can be mapped locally to a VNI at each VTEP, allowing each endpoint to use its own local VLAN configuration while the VNI provides the overlay segment identity. This decoupling is a key scalability benefit of VXLAN: the VNI field is 24 bits, supporting a much larger namespace than traditional VLANs. “Layer 2 frames are encapsulated by the source VTEP.” is also correct. When a frame must traverse the VXLAN overlay, the ingress/source VXLAN Tunnel Endpoint wraps the Layer 2 frame in VXLAN and UDP/IP headers and forwards the resulting packet through the Layer 3 transport network. The receiving VTEP removes the outer encapsulation and delivers the original Ethernet frame into the locally mapped broadcast domain. Juniper describes VTEPs as the devices that encapsulate and decapsulate VXLAN traffic, while

the VXLAN protocol specification defines the VNI as the identifier carried in the VXLAN header. See Juniper's [EVPN-VXLAN overview](#) and [RFC 7348](#).